



---

## EXPLOITING ELEMENTARY LANDSCAPES FOR TSP, VEHICLE ROUTING AND SCHEDULING

Darrell Whitley  
COLORADO STATE UNIVERSITY

---

09/03/2015  
Final Report

DISTRIBUTION A: Distribution approved for public release.

Air Force Research Laboratory  
AF Office Of Scientific Research (AFOSR)/ RTA2  
Arlington, Virginia 22203  
Air Force Materiel Command

<b>REPORT DOCUMENTATION PAGE</b>				Form Approved OMB No. 0704-0188	
<p>The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to the Department of Defense, Executive Service Directorate (0704-0188). Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.</p> <p><b>PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ORGANIZATION.</b></p>					
1. REPORT DATE (DD-MM-YYYY) 31-08-2015		2. REPORT TYPE Final Report		3. DATES COVERED (From - To) From 01-06-2011 To 31-05-2015	
4. TITLE AND SUBTITLE Exploiting Elementary Landscapes for TSP, Vehicle Routing and Scheduling				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER FA9550-11-1-0088	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)  L. Darrell Whitley and Adele E. Howe.				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)  Colorado State University, Fort Collins, CO 80523				8. PERFORMING ORGANIZATION REPORT NUMBER  Air Force Office of Scientific Research	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)  Air Force Office of Scientific Research				10. SPONSOR/MONITOR'S ACRONYM(S)  AFOSR	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT DISTIRBUTION A					
13. SUPPLEMENTARY NOTES  None.					
14. ABSTRACT There are a number of NP-hard optimization problems where the search space can be characterized as an elementary landscape. For these search spaces the evaluation function is an eigenfunction of the Laplacian matrix that describes the neighborhood structure of the search space. Problems such as the Traveling Salesman Problem (TSP) and Graph Coloring are elementary. Problems such as MAX-kSAT are a superposition of k elementary landscapes. This research has exploited statistical and mathematical properties of elementary landscapes to develop new gradient methods for combinatorial optimization problems, particular for TSP and MAXSAT and other k-bounded pseudo-Boolean optimization problems.					
15. SUBJECT TERMS  Combinatorial Optimization, Operations Research, TSP, MAXSAT, k-bounded Pseudo-Boolean Optimization					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON Darrell Whitley
a. REPORT  U	b. ABSTRACT  U	c. THIS PAGE  U			19b. TELEPHONE NUMBER (Include area code) 970-491-5373

## INSTRUCTIONS FOR COMPLETING SF 298

**1. REPORT DATE.** Full publication date, including day, month, if available. Must cite at least the year and be Year 2000 compliant, e.g. 30-06-1998; xx-06-1998; xx-xx-1998.

**2. REPORT TYPE.** State the type of report, such as final, technical, interim, memorandum, master's thesis, progress, quarterly, research, special, group study, etc.

**3. DATES COVERED.** Indicate the time during which the work was performed and the report was written, e.g., Jun 1997 - Jun 1998; 1-10 Jun 1996; May - Nov 1998; Nov 1998.

**4. TITLE.** Enter title and subtitle with volume number and part number, if applicable. On classified documents, enter the title classification in parentheses.

**5a. CONTRACT NUMBER.** Enter all contract numbers as they appear in the report, e.g. F33615-86-C-5169.

**5b. GRANT NUMBER.** Enter all grant numbers as they appear in the report, e.g. AFOSR-82-1234.

**5c. PROGRAM ELEMENT NUMBER.** Enter all program element numbers as they appear in the report, e.g. 61101A.

**5d. PROJECT NUMBER.** Enter all project numbers as they appear in the report, e.g. 1F665702D1257; ILIR.

**5e. TASK NUMBER.** Enter all task numbers as they appear in the report, e.g. 05; RF0330201; T4112.

**5f. WORK UNIT NUMBER.** Enter all work unit numbers as they appear in the report, e.g. 001; AFAPL30480105.

**6. AUTHOR(S).** Enter name(s) of person(s) responsible for writing the report, performing the research, or credited with the content of the report. The form of entry is the last name, first name, middle initial, and additional qualifiers separated by commas, e.g. Smith, Richard, J, Jr.

**7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES).** Self-explanatory.

**8. PERFORMING ORGANIZATION REPORT NUMBER.** Enter all unique alphanumeric report numbers assigned by the performing organization, e.g. BRL-1234; AFWL-TR-85-4017-Vol-21-PT-2.

**9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES).** Enter the name and address of the organization(s) financially responsible for and monitoring the work.

**10. SPONSOR/MONITOR'S ACRONYM(S).** Enter, if available, e.g. BRL, ARDEC, NADC.

**11. SPONSOR/MONITOR'S REPORT NUMBER(S).** Enter report number as assigned by the sponsoring/monitoring agency, if available, e.g. BRL-TR-829; -215.

**12. DISTRIBUTION/AVAILABILITY STATEMENT.** Use agency-mandated availability statements to indicate the public availability or distribution limitations of the report. If additional limitations/ restrictions or special markings are indicated, follow agency authorization procedures, e.g. RD/FRD, PROPIN, ITAR, etc. Include copyright information.

**13. SUPPLEMENTARY NOTES.** Enter information not included elsewhere such as: prepared in cooperation with; translation of; report supersedes; old edition number, etc.

**14. ABSTRACT.** A brief (approximately 200 words) factual summary of the most significant information.

**15. SUBJECT TERMS.** Key words or phrases identifying major concepts in the report.

**16. SECURITY CLASSIFICATION.** Enter security classification in accordance with security classification regulations, e.g. U, C, S, etc. If this form contains classified information, stamp classification level on the top and bottom of this page.

**17. LIMITATION OF ABSTRACT.** This block must be completed to assign a distribution limitation to the abstract. Enter UU (Unclassified Unlimited) or SAR (Same as Report). An entry in this block is necessary if the abstract is to be limited.

Final Report for AFOSR #9550-11-1-0088  
*Exploiting Elementary Landscapes for TSP, Vehicle Routing  
and Scheduling*  
June 1, 2011 to May 31, 2015

Adele Howe     L. Darrell Whitley  
Computer Science Department  
Colorado State University  
Fort Collins, CO 80524  
email: {howe,whitley}@cs.colostate.edu

**Abstract**

There are a number of NP-hard optimization problems where the search space can be characterized as an *elementary landscape*. For these search spaces the evaluation function is an eigenfunction of the Laplacian matrix that describes the neighborhood structure of the search space. Problems such as the Traveling Salesman Problem (TSP), Graph Coloring, the Frequency Assignment Problem, as well as Min-Cut and Max-Cut Graph Partitioning and select simple satisfiability problems all have elementary landscapes. For all elementary landscapes one can compute neighborhood average evaluations without actually evaluating any neighbors. One can also prove that all local optima have an evaluation that is better than the average evaluation over the set of all solutions: there are no arbitrarily poor local optima. And all neighborhoods must contain at least one improving or disimproving move. There are no flat neighborhoods.

We have extended this method for k-bound pseudo-Boolean optimization problems so that we can now compute the exact autocorrelation of the search space as well as the exact statistical moments (mean, variance, skew, kurtosis, ....) over generalized exponentially large neighborhoods at Hamming distance  $d$  for any arbitrary point in the search space. We have also proven that all k-bounded pseudo-Boolean optimization problems, such as MAX-kSAT and NK-landscapes, can be expressed as a superposition of  $k$  elementary landscapes.

We have also exploited key properties of elementary landscapes to develop new search methods that are capable of *tunneling* between local optima, as well as *filtering* local optima. *Tunneling* means that given two local optima, the algorithm is able to construct two (or more) new local optima without any additional search. *Filtering* means that the tunneling operator can reach thousands (even millions) of different local optima, but it automatically selects the best of all of the reachable local optima. For the TSP, tunneling and filtering can be done in  $O(N)$  time; a new algorithm which displays tunneling and filtering behavior has already been shown to find better TSP solutions faster than the Chained-Lin-Kernighan algorithm. For clustered Traveling Salesman Problems and asymmetric TSP our results improve on the state of the art for problems instances too large to solve with Branch and Bound methods (e.g. the Concorde TSP solver).

## Contents

<b>1</b>	<b>Executive Summary</b>	<b>3</b>
1.1	TSP . . . . .	3
1.2	SAT . . . . .	4
1.3	Generalization and Application to Other Problems . . . . .	6
<b>2</b>	<b>Publications</b>	<b>7</b>
2.1	Journal Papers: . . . . .	7
2.2	Conference Papers and Book Chapters: . . . . .	11
2.3	Dissertations and Theses: . . . . .	19

# 1 Executive Summary

We applied the theory of elementary landscapes to compute statistics of NP-hard search spaces, e.g., TSP and MAXSAT, and leveraged the computations to improve search. Among other capabilities, the statistics support more efficient move selection and lookahead in local search.

Interestingly, empirical studies have shown that these statistics are most useful in structured problems. Perhaps it is not surprising that random problems lack structure, and thus summary statistics for random problems are not as informative. But this observation calls into question whether it is a good idea to test search algorithms on benchmarks if those benchmarks are random and devoid of the kinds of structure that is found in real world problems. In our empirical studies, we have focused primarily on applications instances and problems with controlled structure; our philosophy is that algorithms should be designed and evaluated on the kinds of problems that are representative of their likely application.

We were also able to develop new mathematical methods to more automatically identify elementary landscapes. If a combinatorial optimization problem is a superposition of elementary landscapes, algebraic methods can be used to discover a decomposition of the objective function into a sum of subfunctions, where each subfunction is an elementary landscape. Many (but not all) steps of this process can be automated, and we have developed a software tool that assists the researcher in finding an appropriate decomposition [8, 9, 17, 30]. We have also been able to show that some elementary landscapes are decomposable into partial neighborhoods with elementary properties [3, 33].

We have also developed new methods capable of “tunneling” between local optima. This means search can move from a pair of one local optima to another local optimum in a single move; we also show this can be done deterministically at very low cost [13, 14, 20]. This is a radically new way of thinking about the problem of escaping local optima.

Our primary results can be roughly organized by the primary problems of study: TSP, SAT and a set of other problems to which our results can be shown to generalize. These problems are all NP-hard and yet significant progress has been made by us and others on solving large instances and developing explanations of algorithm performance. We have also worked on applications in computer security, specially for resource scheduling [5], security hardening [6] and data anonymization [10, 11]. And we have done some general theory work on comparing algorithms [16] and characterizing problem difficulty [15, 12].

## 1.1 TSP

We have designed a new recombination operator for the TSP called Generalized Partition Crossover (GPX). Recombination algorithms are usually associated with genetic algorithms and evolutionary algorithms. However, our recombination operator can also be embedded in a stochastic local search algorithm; and that is largely how we have used the operator. One novel aspect of our operator is that it is capable of very efficiently “tunneling” between local optima. It can move directly from one pair of local optima to another local optimum in a single move. By integrating the operator in the state-of-the-art multi-trial Lin-Kernighan-Helsgaun 2 (LKH-2) algorithm for the TSP, we have been able to dramatically improve the solutions found for clustered and grid-like TSP instances; this kind of structure is particularly found in circuit layout optimization [27]. These instances are also known to be difficult for LKH-2 alone.

In collaboration with Renato Tinos who visited Colorado State in 2013 from University of Sao Paulo, we have extended to address asymmetric TSPs, producing the Generalized Asym-

metric Partition Crossover (GAPX) [20]. We also applied GAPX to the Dynamic TSP (DTSP) and evaluated the resulting search algorithm with respect to different kinds of dynamically changing environments [19]. For example, one kind of dynamic environment includes the arrival of new tasks (e.g. a new vertex in the graph); another type of dynamic environment results from changing costs (e.g. changing travel times) associated with previously assigned tasks. One reason for studying the asymmetric TSP is that cost metrics such as travel times (as opposed to travel distance) are naturally asymmetric.

Our recombination operators exploit the fact that when we examine two local optima, we can break the solutions down into modular components that are linearly independent. These components can then be recombined to construct new solutions that are linear combinations of the components. Given  $q$  components, our operator is guaranteed to return the best of  $2^q$  new reachable solutions in  $O(n)$  time. The new solutions are guaranteed to be locally optimal in a weak piece-wise linearly additive sense. Empirically we have found that the new solutions are in fact local optima in the majority of cases. This means our search algorithm can move directly from one local optimum to another local optimum, bypassing the exploration of inferior solutions that do not correspond to local optima. It also filters out thousands or even millions of inferior local optima as well [27, 13, 14, 20].

We are continuing to explore the idea that our methods can not only be used to improving heuristic local search methods, but also exact methods. The Concorde branch and bound algorithm for the TSP uses the Chained-Lin-Kernighan algorithm to find a good initial solution; the evaluation of this initial solution is compared to partial solutions associated with different subtrees of the search space. A better initial solution can be used to allow the branch and bound algorithm to prune the search space faster. Adding a recombination operator capable of tunneling greatly improves the Chained-Lin-Kernighan local search algorithm, and thus this might also be used to improve Concorde by providing the branch and bound algorithm with a better initial solution.

## 1.2 SAT

Based on a discrete form of Fourier Analysis called Walsh Analysis, we have constructed closed form calculations to efficiently identify and exploit improving moves for local search algorithms. Additionally, we can compute exact statistics about subregions of the search space [7, 31]. These analytical tools have supported development of efficient hillclimbing algorithms and enhanced exact solvers for the well known Satisfiability problem.

It has long been assumed that “Best Improving” local search is inherently more expensive than “Next Improving” local search. We were able to prove this is not always the case. For MAXSAT problems and other  $k$ -bounded pseudo-Boolean optimization problems, we developed new methods that automatically and analytically identify improving moves, and then track the improving moves so that there is (to a very close approximation of the gradient) no difference in “Best Improving” and “Next Improving” local search. Thus, we have shown that stochastic local search (SLS) can be made dramatically more efficient by identifying best improving moves in constant  $O(1)$  time on average. This technique has been embedded in Best Improving Local Search (BILS) for both NK landscapes [26] and MAXSAT [23, 24] problems. In both cases, we obtained impressive speed-ups in neighborhood evaluations and thus in the number of steps that can be executed in fixed time.

We have developed a new algorithm called WalshMAXSAT that makes use of BILS and the Walsh basis to compute statistics on subspaces of the search space [23]. In our empirical evaluation, WalshMAXSAT finds better local optima five times faster than prior algorithms.

Our study also elucidated an interaction effect between two phases of search in MAXSAT: Stochastic local search (SLS) to a local optimum followed by a perturbation phase. We found that while our fast SLS finds local optima more quickly and even often with better evaluations than the default SLS, the quality of solutions across both phases was sometimes lower; this suggests that best first local search may actually be making plateau traversal (which occurs in the second phase) more difficult for current heuristic solvers. We hypothesize that this is due to the ability of current strategies for traversing plateaus to explore alternative variable settings in some solutions; the locally “better” starting solutions may have more variables set non-optimally when viewed from a global perspective [23].

We have also developed methods to efficiently compute the mean evaluation over Hamming regions. Empirical studies have shown that using the average over a Hamming region as a surrogate objective function can yield superior performance results on neutral landscapes like NKq-landscapes [2]. For MAXSAT, we can identify problem structure by looking for critical variables that appear with greater frequency and that correlate with backbone variables. By exploiting a Fourier polynomial representation of industrial MAXSAT instances, we are able to compute and exploit statistical information about hyperplane slices of the search space associated with these critical variables.

We have shown that this information can be used to improve the performance of algorithms, both SLS-based and exact solvers, for MAXSAT. The hyperplane computation can be leveraged to initialize SLS and thus shorten the number of steps to reach good solutions. The core idea is that the all hyperplanes containing a particular variable votes on the best value based on the expected value of the hyperplane; the value with the most votes becomes the initial setting.

Hyperplane initialization can also be employed as a form of speculative pre-processing for a complete SAT solver. Variable assignments for the 10 most frequently appearing variables are fixed based on the best hyperplane computed from the Walsh analysis. Pruning these variables from the search space can significantly reduce the problem size, allowing the complete solver to more quickly find satisfying assignments or quickly determine that the problem became UNSAT (requiring different assignments of the speculative values). The solver we developed using these ideas performed very well in the Core solver, Sequential, Application SAT track in the SAT 2013 competition: it ranked 3rd in terms of the number of problems solved [21] (however, it was disqualified from the competition due to a problem with the output format which had nothing to do with algorithm performance).

After this initial entry, we studied the interaction between Walsh directed variable assignment and truth preserving preprocessors. From this analysis, we developed a new hybrid solver that selects between a set of possible variable assignments for the speculative phase based on a learned relationship between problem features and the different types of preprocessing. We also submitted a version that used the hyperplane initialization to fix some variables. Our solver for the 2014 competition tied for 5th place in terms of the number of problems solved, and formally in 7th place because our solver was a little slower than two other solvers that solved the same number of MAXSAT problems [22]. We believe this is very good given that some researchers have spent many years of research focused on solving MAXSAT problem. More importantly, our methods represent a new approach. And we can now characterize the statistical structure of the search landscape associated with MAXSAT in far more mathematical detail than has been previously possible.



### 1.3 Generalization and Application to Other Problems

Mathematically well founded gradient descent methods are standard tools in continuous optimization. Such a foundation is lacking in non-linear discrete optimization. The goal of our research is to create a similar mathematical foundation to enable gradient descent methods for discrete combinatorial optimization problems.

We are also generalizing our prior work on TSP by developing new “tunneling” operators for Pseudo-Boolean Optimization (PBO) problems such as MAXSAT, Spin glass problems in theoretical physics and NK-Landscapes from theoretical biology. Specific applications range from hardware verification to RNA folding to search query optimization and task resource scheduling. We have also developed new, more powerful lookahead move selection methods for local search.

We have generalized our algorithm to select improving moves in constant time beyond SAT. Our new methods make it possible to do approximate best improving move local search in  $O(1)$  average time per move for all  $k$ -bounded pseudo-Boolean functions [25, 28]. We have also developed new techniques that can efficiently and automatically identify improving moves at Hamming distance  $r$  for  $k$ -bounded pseudo-Boolean functions [18]. For example, at  $r = 10$ , our algorithms identified all improving moves in the Hamming ball of radius 10 containing  $n^{10}$  points in the search space over binary strings of length  $n$ . We have been able to prove that on average, every single improving move in the Hamming ball is either already known, or is identified in constant time. This means that a local optimum at radius 1 might not be a local optimum at radius 2 or radius 7 or radius 10. In general, each time the radius increases by 1, the number of local optima in the search space is often reduced by an order of magnitude.

Our results can also be used to understanding other types of search algorithms such as evolutionary algorithms and stochastic hill climbing algorithms that work with bit representations. We have collaborated with other researchers to better understand the sampling behavior of these stochastic search algorithms [1] [2] [29] [32]. We have also looked at what makes the class of  $k$  bounded pseudo-Boolean optimization problems difficult [12]

In pseudo-Boolean optimization, the search space is naturally represented as a graph corresponding to a binary hypercube. Simple bit flips move search from one vertex (solution) in the graph to another. Our work has also shown that many of the new computations and methods that we develop for binary representations can be extended and applied to generalized hypercubes with non-binary representations [4]. For example, the search space for Graph Vertex Coloring is naturally represented as a generalized hypercube: instead of a binary representation, the number of colors used to color the graph corresponds to the arity of the generalized hypercube. But the general results that we have shown to hold for the binary hypercube still hold. This means we can still identify the best improving moves in constant  $O(1)$  time on average for problems such as Graph Vertex Coloring. The graph coloring problem also is a prototype for many other scheduling and optimization problems such as time tabling and frequency assignment.

We have continued to exploit the analytical results of the last two or three years and to develop new search methods. In the last year we have been able to exploit tunnels between local optima for  $k$ -bounded pseudo Boolean optimization problems—echoing our earlier results on the Traveling Salesman Problem. Using two locally optimal solutions as “parent” solutions, we have developed a decomposition and recombination operator that deterministically decomposes the parents into  $q$  subsolutions, then reassorted the subsolutions to yield the combination of subsolutions which is guaranteed to have the best possible evaluation compared to all other combinations of the subsolutions. We can also prove that the resulting “offspring” are also

locally optimal in every hyperplane subspace of the search space that contains both of the parents. In practice, this means that the “offspring” of the tunneling recombination are usually (but not always) also locally optimal; if the offspring is not locally optimal, any improving move can only come by changing one of the bit assignments that the parent solutions share in common. This ability to tunnel between local optima has the potential to very much change how iterated local search algorithms are constructed and used.

## 2 Publications

We present a listing of the publications associated with this grant along with the abstracts for the papers.

Our project web site is at: <http://www.cs.colostate.edu/sched/>. From that site, you can access full text of publications and data from the project.

### 2.1 Journal Papers:

1. F. Chicano, A. Sutton, D. Whitley and E. Alba (2015) Fitness Probability Distribution of Bit-Flip Mutation. *Evolutionary Computation*. MIT Press

#### Abstract

Bit-flip mutation is a common mutation operator for evolutionary algorithms applied to optimize functions over binary strings. In this paper, we develop results from the theory of landscapes and Krawtchouk polynomials to exactly compute the probability distribution of fitness values of a binary string undergoing uniform bit-flip mutation. We prove that this probability distribution can be expressed as a polynomial in  $p$ , the probability of flipping each bit. We analyze these polynomials and provide closed-form expressions for an easy linear problem (Onemax), and an NP-hard problem, MAX-SAT. We also discuss a connection of the results with runtime analysis.

2. F. Chicano, D. Whitley and E. Alba (2014) Exact Computation of the Expectation Surfaces for Uniform Crossover along with Bit-flip Mutation. *Theoretical Computer Science*. 545:76-93.

#### Abstract

Uniform crossover and bit-flip mutation are two popular operators used in genetic algorithms to generate new solutions in an iteration of the algorithm when the solutions are represented by binary strings. We use the Walsh decomposition of pseudo-Boolean functions and properties of Krawtchouk matrices to exactly compute the expected value for the fitness of a child generated by uniform crossover followed by bit-flip mutation from two parent solutions. We prove that this expectation is a polynomial in  $p$ , the probability of selecting the best-parent bit in the crossover, and  $q$ , the probability of flipping a bit in the mutation. We provide efficient algorithms to compute this polynomial for Onemax and MAX-SAT problems, but the results also hold for other problems such as NK-Landscapes. We also analyze the features of the expectation surfaces.

3. D. Whitley, A. Sutton, G. Ochoa and F. Chicano (2014) The Component Model for Elementary Landscapes and Partial Neighborhoods. *Theoretical Computer Science*. Vol 545:59-75.

#### Abstract

Local search algorithms exploit moves on an adjacency graph of the search space. An "elementary landscape" exists if the objective function  $f$  is an eigenfunction of the Laplacian of the graph induced by the neighborhood operator; this allows various statistics about the neighborhood to be computed in closed form. A new component based model makes it relatively simple to prove that certain types of landscapes are elementary. The traveling salesperson problem, weighted graph (vertex) coloring and the minimum graph bisection problem yield elementary landscapes under commonly used local search operators. The component model is then used to efficiently compute the mean objective function value over partial neighborhoods for these same problems. For a traveling salesperson problem over  $n$  cities, the 2-opt neighborhood can be decomposed into  $n/2-1$  partial neighborhoods. For graph coloring and the minimum graph bisection problem, partial neighborhoods can be used to focus search on those moves that are capable of producing a solution with a strictly improving objective function value.

4. A. Sutton, F. Chicano, D. Whitley (2013) Fitness function distributions over generalized search neighborhoods in the  $q$ -ary hypercube. *Evolutionary Computation* 21(4):561-590.

#### Abstract

The frequency distribution of a fitness function over regions of its domain is an important quantity for understanding the behavior of algorithms that employ randomized sampling to search the function. In general, exactly characterizing this distribution is at least as hard as the search problem, since the solutions typically live in the tails of the distribution. However, in some cases it is possible to efficiently retrieve a collection of quantities (called moments) that describe the distribution. In this paper, we consider functions of bounded epistasis that are defined over length- $n$  strings from a finite alphabet of cardinality  $q$ . Many problems in combinatorial optimization can be specified as search problems over functions of this type. Employing Fourier analysis of functions over finite groups, we derive an efficient method for computing the exact moments of the frequency distribution of fitness functions over Hamming regions of the  $q$ -ary hypercube. We then use this approach to derive equations that describe the expected fitness of the offspring of any point undergoing uniform mutation. The results we present provide insight into the statistical structure of the fitness function for a number of combinatorial problems. For the graph coloring problem, we apply our results to efficiently compute the average number of constraint violations that lie within a certain number of steps of any coloring. We derive an expression for the mutation rate that maximizes the expected fitness of an offspring at each fitness level. We also apply the results to the slightly more complex frequency assignment problem, a relevant application in the domain of the telecommunications industry. As with the graph coloring problem, we provide formulas for the average value of the fitness function in Hamming regions around a solution and the expectation-optimal mutation rate.

5. Rinku Dewri, Indrajit Ray, Indrakshi Ray, D. Whitley (2012) Utility Driven Optimization of Real Time Data Broadcast Schedules. *Applied Soft Computing*. 12(7):1832-1846.

#### **Abstract**

Data dissemination in wireless environments is often accomplished by on-demand broadcasting. The time critical nature of the data requests plays an important role in scheduling these broadcasts. Most research in on-demand broadcast scheduling has focused on the timely servicing of requests so as to minimize the number of missed deadlines. However, there exists many environments where the utility of the received data is an equally important criterion as its timeliness. Missing the deadline may reduce the utility of the data but does not necessarily make it zero. In this work, we address the problem of scheduling real time data broadcasts with such soft deadlines. We investigate search based optimization techniques to develop broadcast schedulers that make explicit attempts to maximize the utility of data requests as well as service as many requests as possible within an acceptable time limit. Our analysis shows that heuristic driven methods for such problems can be improved by hybridizing them with local search algorithms. We further investigate the option of employing a dynamic optimization technique to facilitate utility gain, thereby eliminating the requirement of a heuristic in the process. An evolution strategy based stochastic hill-climber is investigated in this context.

6. Rinku Dewri, Indrajit Ray, Nayot Poolsappasit, D. Whitley (2012) Optimal Security Hardening on Attack Tree Models of Networks: A Cost-Benefit Analysis. *International Journal of Information Security*. 11(3):167-188.

#### **Abstract**

Researchers have previously looked into the problem of determining whether a given set of security hardening measures can effectively make a networked system secure. However, system administrators are often faced with a more challenging problem since they have to work within a fixed budget which may be less than the minimum cost of system hardening. An attacker, on the other hand, explores alternative attack scenarios to inflict the maximum damage possible when the security controls are in place, very often rendering the optimality of the controls invalid. In this work, we develop a systematic approach to perform a cost-benefit analysis on the problem of optimal security hardening under such conditions. Using evolutionary paradigms such as multi-objective optimization and competitive co-evolution, we model the attacker-defender interaction as an 'arms race', and explore how security controls can be placed in a network to induce a maximum return on investment.

7. A. Sutton, D. Whitley, A. Howe (2012) Computing the moments of k-bounded pseudo-Boolean functions over Hamming spheres of arbitrary radius in polynomial time. *Theoretical Computer Science*, 425:58-74.

#### **Abstract**

We show that given a k-bounded pseudo-Boolean function  $f$ , one can always compute the  $c$ th moment of  $f$  over regions of arbitrary radius in Hamming space in polynomial time using algebraic information from the adjacency structure (where  $k$  and  $c$  are constants). This result has implications for evolutionary algorithms and local search algorithms because information about promising regions of the search space can be efficiently retrieved, even if the cardinality of the region is exponential in the problem size. Finally, we use our results to introduce a method of efficiently calculating the expected fitness of mutations for evolutionary algorithms.

8. F. Chicano, D. Whitley, E. Alba (2011) A Methodology to Find the Elementary Landscape Decomposition of Combinatorial Optimization Problems. *Evolutionary Computation*, 19(4):597-637.

#### Abstract

A small number of combinatorial optimization problems have search spaces that correspond to elementary landscapes, where the objective function  $f$  is an eigenfunction of the Laplacian that describes the neighborhood structure of the search space. Many problems are not elementary; however, the objective function of a combinatorial optimization problem can always be expressed as a superposition of multiple elementary landscapes if the underlying neighborhood used is symmetric. This paper presents theoretical results that provide the foundation for algebraic methods that can be used to decompose the objective function of an arbitrary combinatorial optimization problem into a sum of subfunctions, where each subfunction is an elementary landscape. Many steps of this process can be automated, and indeed a software tool could be developed that assists the researcher in finding a landscape decomposition. This methodology is then used to show that the subset sum problem is a superposition of two elementary landscapes, and to show that the quadratic assignment problem is a superposition of three elementary landscapes.

9. F. Chicano, D. Whitley, E. Alba, F Luna (2011) Elementary Landscape Decomposition of the Frequency Assignment Problem. *Theoretical Computer Science* 412(43,7):6002-6019.

#### Abstract

The Frequency Assignment Problem (FAP) is an important problem that arises in the design of radio networks, when a channel has to be assigned to each transceiver of the network. This problem is a generalization of the graph coloring problem. In this paper we study a general version of the FAP that can include adjacent frequency constraints. Using concepts from landscapes? theory, we prove that this general FAP can be expressed as a sum of two elementary landscapes. Further analysis also shows that some subclasses of the problem correspond to a single elementary landscape. This allows us to compute the kind of neighborhood information that is normally associated with elementary landscapes. We also provide a closed form formula for computing the autocorrelation coefficient for the general FAP, which can be useful as an a priori indicator of the performance of a local search method.

10. Rinku Dewri, Indrajit Ray, Indrakshi Ray, D. Whitley (2011) Exploring Privacy versus Data Quality Trade-offs in Anonymization using Multi-objective Optimization. *Journal of Computer Security* 19(5). (May 2011).

#### Abstract

Data anonymization techniques have received extensive attention in the privacy research community over the past several years. Various models of privacy preservation have been proposed:  $k$ -anonymity,  $\ell$ -diversity and  $t$ -closeness, to name a few. An oft-cited drawback of these models is that there is considerable loss in data quality arising from the use of generalization and suppression techniques. Optimization attempts in this context have so far focused on maximizing the data utility for a pre-specified level of privacy. To determine if better privacy levels are obtainable with the same level of data utility, majority of the existing formulations require exhaustive analysis. Further, the data publisher's perspective is often missed in the process. The publisher wishes to maintain a given level of data utility since the data utility is the revenue earner and

then maximize the level of privacy within acceptable limits. In this paper, we explore this privacy versus data quality trade-off as a multi-objective optimization problem. Our goal is to provide substantial information to a data publisher about the trade-offs available between the privacy level and the information content of an anonymized data set.

11. Rinku Dewri, Indrakshi Ray, Indrajit Ray, D. Whitley (2011) k-Anonymization in the Presence of Publisher Preferences. *Transactions on Knowledge and Data Engineering*. (May 2011).

### Abstract

Privacy constraints are typically enforced on shared data that contain sensitive personal attributes. However, owing to its adverse effect on the utility of the data, information loss must be minimized while sanitizing the data. Existing methods for this purpose modify the data only to the extent necessary to satisfy the privacy constraints, thereby asserting that the information loss has been minimized. However, given the subjective nature of information loss, it is often difficult to justify such an assertion. In this paper, we propose an interactive procedure to generate a data generalization scheme that optimally meets the preferences of the data publisher. A data publisher guides the sanitization process by specifying aspirations in terms of desired achievement levels in the objectives. A reference direction based methodology is used to investigate neighborhood solutions if the generated scheme is not acceptable. This approach draws its power from the constructive input received from the publisher about the suitability of a solution before finding a new one.

## 2.2 Conference Papers and Book Chapters:

12. D. Whitley (2015) Mk Landscapes, NK Landscapes and MAX-kSAT: A Proof that the Only Challenging Problems are Deceptive. *Genetic and Evolutionary Computation Conference*, GECCO-2015. ACM Press.

### Abstract

This paper investigates *Gray Box Optimization* for pseudo-Boolean optimization problems composed of  $M$  subfunctions, where each subfunction accepts at most  $k$  variables. We will refer to these as Mk Landscapes. In Gray Box optimization, the optimizer is given access to the set of  $M$  subfunctions. If the set of subfunctions is  $k$ -bounded and separable, the Gray Box optimizer is guaranteed to return the global optimum with 1 evaluation. A problem is said to be *order  $k$  deceptive* if the average values of hyperplanes over combinations of  $k$  variables cannot be used to infer a globally optimal solution. Hyperplane averages are always efficiently computable for Mk Landscapes. If a problem is not deceptive, the Gray Box optimizer also returns the global optimum after 1 evaluation. Finally, these concepts are used to understand the nonlinearity of problems in the complexity class P, such as Adjacent NK Landscapes. These ideas are also used to understand the problem structure of NP Hard problems such as MAX-kSAT and general Mk Landscapes. In general, NP Hard problems are profoundly deceptive.

13. G. Ochoa, F. Chicano, R. Tinos and D. Whitley (2015) Tunnelling Crossover Networks. *Genetic and Evolutionary Computation Conference, GECCO-2015*. ACM Press.

#### Abstract

Local optima networks are a recent model of fitness landscapes. They compress the landscape by representing local optima as nodes, and search transitions among them as edges. Previous local optima networks considered transitions based on mutation; this study looks instead at transitions based on deterministic recombination. We define and analyze networks based on the recently proposed partition crossover for  $k$ -bounded pseudo-Boolean functions, using  $NK_q$  landscapes as a case study. Partition crossover was initially proposed for the traveling salesman problem, where it was found to “tunnel” between local optima, i.e. jump from local optimum to local optimum. Our network analysis shows that this is also the case for  $NK$  landscapes: local optima are densely connected via partition crossover. We found marked differences between the adjacent and random interaction  $NK$  models. Surprisingly, with the random model, instances have a lower number of local optima on average, but their networks are more sparse and decompose into several clusters. There is also large variability in the size and pattern of connectivity of instances coming from the same landscape parameter values. These network features offer new insight informing why some instances are harder to solve than others.

14. R. Tinos, D. Whitley and F. Chicano (2015) Partition Crossover for Pseudo-Boolean Optimization. *Foundations of Genetic Algorithms, FOGA-2015*. ACM Press.

#### Abstract

A generalized partition recombination operator is introduced for  $k$ -bounded pseudo-Boolean functions such as  $NK$  landscapes and MAX- $k$ SAT. Under partition crossover, the evaluation of the offspring can be directly obtained from partial evaluations in the parents. Partition crossover explores the interaction graph of the pseudo-Boolean functions to partition the variables of the solution vector. If the differing variables in the two parents can be partitioned into  $q$  non-interacting sets, partition crossover can be used to find the best of  $2^q$  possible offspring. Experimental results show the efficiency of the proposed crossover when used in combination with a hybrid genetic algorithm.

15. D. Whitley and J. Rowe (2014) Single Funnel and Multi-Funnel Landscapes and Subthreshold Seeking Behavior. In : Y. Borenstein and A. Moraglio, eds., *Theory and Principled Methods for the Design of Metaheuristics*, Springer. pp: 63-84.

#### Abstract

Algorithms for parameter optimization display subthreshold-seeking behavior when the majority of the points that the algorithm samples have an evaluation less than some target threshold. Subthreshold-seeking algorithms avoid the curse of the general and Sharpened No Free Lunch theorems in the sense that they are better than random enumeration on a specific (but general) family of functions. In order for subthreshold-seeking search to be possible, most of the solutions that are below threshold must be localized in one or more regions of the search space. Functions with search landscapes that can be characterized as single-funnel or multi-funnel landscapes have this localized property. We first analyze a simple “Subthreshold-Seeker” algorithm. Further theoretical analysis details conditions that would allow a Hamming neighborhood local search algorithm using a Gray or binary representation to display subthreshold-seeking behavior. A very simple modification to local search is proposed that improves its subthreshold-seeking behavior.

16. D. Whitley (2014) Sharpened and Focused No Free Lunch Theorem and Complexity Theory In, *Search Methodologies, 2nd Edition* E. Burke and G. Kendall, eds., Springer.

#### Abstract

This tutorial reviews basic concepts in complexity theory, as well as various No Free Lunch results and how these results relate to computational complexity. The tutorial explains basic concepts in an informal fashion that illuminates key concepts. ?No Free Lunch? theorems for search can be summarized by the following result: For all possible performance measures, no search algorithm is better than another when its performance is averaged over all possible discrete functions. Note that No Free Lunch is often referred to simply as NFL within the heuristic search community (despite copyrights and trademarks held by the National Football League). Two more recent variants of NFL, the Sharpened NFL, and the Focused NFL are also reviewed. There has been a significant amount of confusion in the literature about the meaning of No Free Lunch, and differences between Sharpened NFL and Focused NFL have not been well understood in the literature. This tutorial attempts to resolve some of this confusion. The reader familiar with basic complexity theory might wish to skip to Sect. 16.3 on No Free Lunch. Section 16.4 explains new results based on the distinction between Sharpened and Focused No Free Lunch.

17. F. Chicano and D. Whitley (2014) Elementary Landscape Decomposition of the Hamiltonian Path Optimization Problem. *EvoCop/EvoStar 2014*.

#### Abstract

There exist local search landscapes where the evaluation function is an eigenfunction of the graph Laplacian that corresponds to the neighborhood structure of the search space. Problems that display this structure are called “Elementary Landscapes” and they have a number of special mathematical properties. In this paper, the Hamiltonian Path problem is shown to be elementary.

18. F. Chicano, D. Whitley and A. Sutton (2014) Efficient Identification of Improving Moves in a Ball for Pseudo-Boolean Problems. *Genetic and Evolutionary Computation Conference*, GECCO-2014. ACM Press.

#### Abstract

Hill climbing algorithms are at the core of many approaches to solve optimization problems. Such algorithms usually require the complete enumeration of a neighborhood of the current solution. In the case of problems defined over binary strings of length  $n$ , we define the  $r$ -ball neighborhood as the set of solutions at Hamming distance  $r$  or less from the current solution. For  $r \ll n$  this neighborhood contains  $\Theta(n^r)$  solutions. In this paper efficient methods are introduced to locate improving moves in the  $r$ -ball neighborhood for problems that can be written as a sum of a linear number of subfunctions depending on a bounded number of variables. NK-landscapes and MAX-kSAT are examples of these problems. If the number of subfunctions depending on any given variable is also bounded, then we prove that the method can explore the neighborhood in constant time, despite the fact that the number of solutions in the neighborhood is polynomial in  $n$ . We develop a hill climber based on our exploration method and we analyze its efficiency and efficacy using experiments with NKq-landscapes instances.



19. R. Tinos, D. Whitley, and A. Howe (2014) Using of Explicit Memory for the Dynamic Traveling Salesman Problem. *Genetic and Evolutionary Computation Conference, GECCO-2014*. ACM Press.

#### Abstract

In the dynamic traveling salesman problem (DTSP), the weights and vertices of the graph representing the TSP are allowed to change during the optimization. This work first discusses some issues related to the use of evolutionary algorithms in the DTSP. From simulations of efficient algorithms used for the static TSP applied to optimize solutions in the environments between two random changes in the DTSP, we observe that only some edges, from all possible edges, are generally inserted in and removed from the best solutions after the changes, indicating a possible beneficial use of memory approaches, usually employed in cyclic dynamic environments. We propose a memory approach and a hybrid approach that combines our memory approach with Elitism-based immigrants GA (EIGA). We compare these two approaches to four existing approaches and show that memory approaches can be beneficial for the DTSP with random changes.

20. R. Tinos, D. Whitley, and G. Ochoa (2014) Generalized Asymmetric Partition Crossover for the Asymmetric TSP. *Genetic and Evolutionary Computation Conference, GECCO-2014*. ACM Press.

#### Abstract

The Generalized Partition Crossover (GPX) constructs new solutions for the symmetric Traveling Salesman Problem (TSP) by finding recombining partitions with two cutting points in the graph composed by the union of two parent solutions. If there are  $k$  recombining partitions in the union graph,  $2^k - 2$  solutions simultaneously exploited by GPX. A new Generalized Asymmetric Partition Crossover (GAPX) is introduced that can find more partitions and can also find feasible partitions for the Asymmetric Traveling Salesman Problem. The most important characteristics of the new operator are that it is capable to find recombining partitions that break vertices of degree 4, and it is able to find partition with multiple entry and exit points. Experimental results show that GAPX is capable to improve the quality of solutions generated by the LKH heuristic.

21. D. Hains, A. Howe and D. Whitley (2013) MiniSAT with Speculative Preprocessing, *Proceedings of SAT COMPETITION 2013: Solver and Benchmark Descriptions*, pp. 50-51.

#### Abstract

We present a new preprocessing technique utilizing subspace averages to perform a reduction by heuristically fixing truth assignments. We use the Walsh transform to efficiently compute the average evaluation of solutions in subspaces of the search space that we refer to as hyperplanes. A hyperplane contains all solutions that have the same truth assignments over some subset of the  $n$  variables in the given formula. The formula is reduced by fixing the values from the hyperplane with the best average and reduce the original formula by eliminating clauses satisfied by the variables with consistent truth assignments across all solutions in the hyperplane. We then run the MiniSAT complete solver on this reduced instance.

22. W. Chen, A. Howe and D. Whitley (2014) MiniSAT with Classification-based Preprocessing, *Proceedings of SAT COMPETITION 2014: Solver and Benchmark Descriptions*, pp. 41-42.

#### Abstract

We present a classification-based approach to selecting preprocessors of CNF formulas for the complete solver MiniSAT. Three different preprocessors are considered prior to running MiniSAT. To obtain training data for classification, each preprocessor is run, followed by running MiniSAT on its resulting CNF, on instances from the last competition. On each instance, the preprocessor leading to the smallest solving time is considered the “best” and used as the classification for the instance. A decision tree is trained to select the best preprocessor according to features computed from the CNF formulas. The decision tree can therefore be used to predict the best preprocessor for new problem instances.

23. D. Whitley, A. Howe and D. Hains (2013) Greedy or Not? Best Improving versus First Improving Stochastic Local Search for MAXSAT. *AAAI 2013*.

#### Abstract

Stochastic local search (SLS) is the dominant paradigm for incomplete SAT and MAXSAT solvers. Early studies on small 3SAT instances found that the use of “best improving” moves did not improve search compared to using an arbitrary “first improving” move. Yet SLS algorithms continue to use best improving moves. We revisit this issue by studying very large random and industrial MAXSAT problems. Because locating best improving moves is more expensive than first improving moves, we designed an “approximate best” improving move algorithm and prove that it is as efficient as first improving move SLS. For industrial problems the first local optima found using best improving moves are statistically significantly better than local optima found using first improving moves. However, this advantage reverses as search continues and algorithms must explore equal moves on plateaus. This reversal appears to be associated with critical variables that are in many clauses and that also yield large improving moves.

24. D. Hains, D. Whitley, A. Howe and W. Chen (2013) Hyperplane Initialized Local Search for MAXSAT. *Genetic and Evolutionary Computation Conference, GECCO-2013*. ACM Press.

#### Abstract

By converting the MAXSAT problem to Walsh polynomials, we can efficiently and exactly compute the hyperplane averages of fixed order  $k$ . We use this fact to construct initial solutions based on variable configurations that maximize the sampling of hyperplanes with good average evaluations. The Walsh coefficients can also be used to implement a constant time neighborhood update which is integral to a fast next descent local search for MAXSAT (and for all bounded pseudo-Boolean optimization problems.) We evaluate the effect of initializing local search with hyperplane averages on both the first local optima found by the search and the final solutions found after a fixed number of bit flips. Hyperplane initialization not only provides better evaluations, but also finds local optima closer to the globally optimal solution in fewer bit flips than search initialized with random solutions. A next descent search initialized with hyperplane averages is able to outperform several state-of-the art stochastic local search algorithms on both random and industrial instances of MAXSAT.

25. W. Chen, D. Whitley, A. Howe and D. Hains (2013) Second Order Derivatives for NK-Landscapes. *Genetic and Evolutionary Computation Conference, GECCO-2013*. ACM Press.

#### Abstract

Local search methods based on explicit neighborhood enumeration requires at least  $O(n)$  time to identify all possible improving moves. For  $k$ -bounded pseudo-Boolean optimization problems, recent approaches have achieved  $O(k^2 * 2^k)$  runtime cost, where  $n$  is the number of variables and  $k$  is the number of variables per subfunction. Even though the bound is independent of  $n$ , the complexity per move is still exponential in  $k$ . In this paper, we propose a second order partial derivatives-based approach that executes first-improvement local search where the runtime cost per move is time polynomial in  $k$  and independent of  $n$ . This method is applied to NK-landscapes, where larger values of  $k$  may be of particular interest.

26. D. Whitley, W. Chen and A. Howe (2012) An Empirical Evaluation of  $O(1)$  Steepest Descent for NK-Landscapes. *Parallel Problem Solving from Nature, PPSN 12*, Springer.

#### Abstract

New methods make it possible to do approximate steepest descent in  $O(1)$  time per move for  $k$ -bounded pseudo-Boolean functions using stochastic local search. It is also possible to use the average fitness over the Hamming distance 2 neighborhood as a surrogate fitness function and still retain the  $O(1)$  time per move. These are average complexity results. In light of these new results, we examine three factors that can influence both the computational cost and the effectiveness of stochastic local search: 1) Fitness function:  $f(x)$  or a surrogate; 2) Local optimum escape method: hard random or soft restarts; 3) Descent strategy: next or steepest. We empirically assess these factors in a study of local search for solving NK-landscape problems.

27. D. Hains, D. Whitley and A. Howe (2012) Improving Lin-Kernighan-Helsgaun with Crossover on Clustered Instances of the TSP. *Parallel Problem Solving from Nature, PPSN 12*, Springer.

#### Abstract

Multi-trial Lin-Kernighan-Helsgaun 2 (LKH-2) is widely considered to be the best Iterated Local Search heuristic for the Traveling Salesman Problem (TSP) and has found the best-known solutions to a large number of benchmark problems. Although LKH-2 performs exceptionally well on most instances, it is known to have difficulty on clustered instances of the TSP. Generalized Partition Crossover (GPX) is a crossover operator for the TSP that efficiently constructs new solutions by partitioning a graph constructed from the union of two solutions. We show that GPX is especially well-suited for clustered instances and evaluate its ability to improve solutions found by LKH-2. We present two methods of combining GPX with multi-trial LKH-2. We find that combining GPX with LKH-2 dramatically improves the evaluation of solutions found by LKH-2 alone on clustered instances with sizes ranging from 3,000 to 30,000 cities.

28. D. Whitley and W. Chen (2012) Constant Time Steepest Descent Local Search with Lookahead for NK-Landscapes and MAX-kSAT. *Genetic and Evolutionary Computation Conference*, GECCO-2012. ACM Press.

#### Abstract

A modified form of steepest descent local search is proposed that displays an average complexity of  $O(1)$  time per move for NK-Landscape and MAX-kSAT problems. The algorithm uses a Walsh decomposition to identify improving moves. In addition, it is possible to compute a Hamming distance 2 statistical lookahead: if  $x$  is the current solution and  $y$  is a neighbor of  $x$ , it is possible to compute the average evaluation of the neighbors of  $y$ . The average over the Hamming distance 2 neighborhood can be used as a surrogate evaluation function to replace  $f$ . The same modified steepest descent can be executed in  $O(1)$  time using the Hamming distance 2 neighborhood average as the fitness function. In practice, the modifications needed to prove  $O(1)$  complexity can be relaxed with little or no impact on runtime performance. Finally, steepest descent local search over the mean of the Hamming distance 2 neighborhood yields superior results compared to using the standard evaluation function for certain types of NK-Landscape problems.

29. F. Chicano, D. Whitley and E. Alba (2012) Exact Computation of the Expectation Curves for Uniform Crossover *Genetic and Evolutionary Computation Conference*, GECCO-2012. ACM Press.

#### Abstract

Bit-flip mutation is a common operation when a genetic algorithm is applied to solve a problem with binary representation. We use in this paper some results of landscapes theory and Krawtchouk polynomials to exactly compute the expected value of the fitness of a mutated solution. We prove that this expectation is a polynomial in  $p$ , the probability of flipping a single bit. We analyze these polynomials and propose some applications of the obtained theoretical results.

30. D. Whitley and F. Chicano (2012) Quasi-Elementary Landscapes and Superpositions of Elementary Landscapes. *Learning and Intelligent Optimization Conference*, LION-2012.

#### Abstract

There exist local search landscapes where the evaluation function is an eigenfunction of the graph Laplacian that corresponds to the neighborhood structure of the search space. Problems that display this structure are called “Elementary Landscapes” and they have a number of special mathematical properties. The term “Quasi-elementary landscapes” is introduced to describe landscapes that are “almost” elementary; in quasi-elementary landscapes there exists some efficiently computed “correction” that captures those parts of the neighborhood structure that deviate from the normal structure found in elementary landscapes. The “shift” operator, as well as the “3-opt” operator for the Traveling Salesman Problem landscapes induce quasi-elementary landscapes. A local search neighborhood for the Maximal Clique problem is also quasi-elementary. Finally, we show that landscapes which are a superposition of elementary landscapes can be quasi-elementary in structure.

31. A. Sutton, D. Whitley and A.E. Howe (2011) Approximating the Distribution of Fitness over Hamming Regions. *Foundations of Genetic Algorithms-2011* ACM Press.

**Abstract**

The distribution of fitness values across a set of states sharply influences the dynamics of evolutionary processes and heuristic search in combinatorial optimization. In this paper we present a method for approximating the distribution of fitness values over Hamming regions by solving a linear programming problem that incorporates low order moments of the target function. These moments can be retrieved in polynomial time for select problems such as MAX- $k$ -SAT using Walsh analysis. The method is applicable to any real function on binary strings that is epistatically bounded and discrete with asymptotic bounds on the cardinality of its codomain.

We perform several studies on the ONE-MAX and MAX- $k$ -SAT domains to assess the accuracy of the approximation and its dependence on various factors. We show that the approximation can accurately predict the number of states within a Hamming region that have an improving fitness value.

32. A. Sutton, D. Whitley and A.E. Howe (2011) Mutation Rates of the (1+1)EA on Pseudo-Boolean Functions of Bounded Epistasis. *Genetic and Evolutionary Computation Conference*, GECCO-2011. ACM Press.

**Abstract**

When the epistasis of the fitness function is bounded by a constant, we show that the expected fitness of an offspring of the (1+1)-EA can be efficiently computed for any point. Moreover, we show that, for any point, it is always possible to efficiently retrieve the “best” mutation rate at that point in the sense that the expected fitness of the resulting offspring is maximized.

On linear functions, it has been shown that a mutation rate of  $1/n$  is provably optimal. On functions where epistasis is bounded by a constant  $k$ , we show that for sufficiently high fitness, the commonly used mutation rate of  $1/n$  is also best, at least in terms of maximizing the expected fitness of the offspring. However, we find for certain ranges of the fitness function, a better mutation rate can be considerably higher, and can be found by solving for the real roots of a degree- $k$  polynomial whose coefficients contain the nonzero Walsh coefficients of the fitness function. Simulation results on maximum  $k$ -satisfiability problems and NK-landscapes show that this expectation-maximized mutation rate can cause significant gains early in search.

33. D. Whitley and G. Ochoa (2011) Partial Neighborhoods of the Traveling Salesman Problem. *Genetic and Evolutionary Computation Conference*, GECCO-2011. ACM Press.

**Abstract**

The Traveling Salesman Problem (TSP) is known to display an elementary landscape under all  $k$ -opt move operators. Previous work has also shown that partial neighborhoods may exist that retain some properties characteristic of elementary landscapes. For a tour of  $n$  cities, we show that the 2-opt neighborhood can be decomposed into  $\lfloor n/2 - 1 \rfloor$  partial neighborhoods. While this paper focuses on the TSP, it also introduces a more formal treatment of partial neighborhoods which applies to all elementary landscapes. Tracking partial neighborhood averages in elementary landscapes requires partitioning the cost matrix. After every move in the search space, the relevant partitions must be updated. However, just as the evaluation function allows a partial update for the TSP, there also exists a partial update for the cost matrix partitions. By only looking at a subset of the partial neighborhoods we can further reduce the cost of updating the cost matrix partitions.

## 2.3 Dissertations and Theses:

**Doug Hains, Fall 2013. Ph.D. Dissertation: Structure in Combinatorial Optimization and its Effect on Heuristic Performance, Colorado State University.** The goal in combinatorial optimization is to find a good solution among a finite set of solutions. In many combinatorial problems, the set of solutions scales at an exponential or greater rate with the instance size. The maximum boolean satisfiability (MAX-SAT) is one such problem that has many important theoretical and practical applications. Due to the exponential growth of the search space, sufficiently large instances of MAX-SAT are intractable for complete solvers. Incomplete solvers, such as stochastic local search (SLS) algorithms are necessary to find solutions in these cases. Many SLS algorithms for MAXSAT have been developed on randomly generated benchmarks using a uniform distribution. As a result, SLS algorithms for MAX-SAT perform exceptionally well on uniform random instances. However, instances from real-world applications of MAX-SAT have a structure that is not captured in expectation by uniform random problems. The same SLS algorithms that perform well on uniform instances have a drastic drop in performance on structured instances.

To better understand the performance drop on structured instances, we examine three characteristics commonly found in real-world applications of MAX-SAT: a power-law distribution of variables, clause lengths following a power-law distribution, and a community structure similar to that found in small-world models. We find that those instances with a community structure and clause lengths following a power-law distribution have a significantly more rugged search space and larger backbones than uniform random instances. These search space properties make it more difficult for SLS algorithms to find good solutions and in part explains the performance drop on industrial instances.

In light of these findings, we examine two ways of improving the performance of SLS algorithms on industrial instances. First, we present a method of tractably computing the average evaluation of solutions in a subspace that we call a hyperplane. These averages can be used to estimate the correct setting of the backbone variables, with as high as 90% accuracy on industrial-like instances. By initializing SLS algorithms with these solutions, the search is able to find significantly better solutions than using standard initialization methods. Second, we reexamine the trade-offs between first and best improving search. We find that in many cases, the evaluation of solutions found by SLS algorithms using first improving search are no worse, and sometimes better, than those found by best improving. First improving search is significantly faster; using first improving search with AdaptG2WSAT, a state-of-the-art SLS algorithm for MAX-SAT, gives us more than a 1,000x speedup on large industrial instances.

Finally, we use our hyperplane averages to improve the performance of complete solvers of the satisfiability problem (SAT), the decision version of MAX-SAT. We use the averages to heuristically select a promising hyperplane and perform a reduction of the original problem based on the chosen hyperplane. This significantly reduces the size of the space that must be searched by the complete solver. Using hyperplane reduction as a preprocessing step, a standard complete SAT solver is able to outperform many state-of-the-art complete solvers. Our contributions have advanced the understanding of structured instances and the performance of both SLS algorithms and complete solvers on these instances. We also hope that this work will serve as a foundation for developing better heuristics and complete solvers for real-world applications of SAT and MAX-SAT.

**Wenxiang Chen, Fall 2013. M.S. Thesis: A Step Toward Constant Time Local Search For Optimizing Pseudo Boolean Functions, Colorado State University.**

Pseudo Boolean Functions (PBFs) are the objective functions for a wide class of hard optimization problems, such as MAX-SAT and MAX-CUT. Since these problems are NPHard, researchers and practitioners rely on incomplete solvers, such as Stochastic Local Search (SLS), for large problems. Best-Improvement Local Search (BILS) is a common form of SLS, which always takes the move yielding the highest improvement in the objective function. Generally, the more runtime SLS is given, the better solution can be obtained. This thesis aims at algorithmically accelerating SLS for PBFs using Walsh Analysis.

The contributions of this thesis are threefold. First, a general approach for executing an approximate best-improvement move in constant time on average using Walsh analysis, Walsh-LS”, is described. Conventional BILS typically requires examining all  $n$  neighbors to decide which move to take, given the number of variables is  $n$ . With Walsh analysis, however, we can determine which neighbors need to be checked. As long as the objective function is epistatically bounded by a constant  $k$  ( $k$  is the number of variables per subfunctions), the number of neighbors that need to be checked is constant regardless of problem size. An impressive speedup of runtime (up to 449X) is observed in our empirical studies.

Second, in the context of Walsh-LS, we empirically study two key components of SLS from the perspectives of both efficiency and effectiveness: 1) Local optimum escape method: hard random or soft restarts; 2) Local search strategy: first-improvement or best-improvement.

Lastly, on average we can perform approximate BILS using the mean over a Hamming region of arbitrary radius as a surrogate objective function. Even though the number of points is exponential in the radius of the Hamming region, BILS using the mean value of points in the Hamming region as a surrogate objective function can still take each move in time independent of  $n$  on average. According to our empirical studies, using the average over a Hamming region as a surrogate objective function can yield superior performance results on neutral landscapes like NKq-landscapes.

1.

**1. Report Type**

Final Report

**Primary Contact E-mail****Contact email if there is a problem with the report.**

whitley@cs.colostate.edu

**Primary Contact Phone Number****Contact phone number if there is a problem with the report**

970-491-5373

**Organization / Institution name**

Colorado State University

**Grant/Contract Title****The full title of the funded effort.**

Exploiting Elementary Landscapes for TSP, Vehicle Routing and Scheduling

**Grant/Contract Number****AFOSR assigned control number. It must begin with "FA9550" or "F49620" or "FA2386".**

FA9550-11-1-0088

**Principal Investigator Name****The full name of the principal investigator on the grant or contract.**

Darrell Whitley and Adele Howe

**Program Manager****The AFOSR Program Manager currently assigned to the award**

Fariba Fahroo

**Reporting Period Start Date**

06/01/2011

**Reporting Period End Date**

05/31/2015

**Abstract**

There are a number of NP-hard optimization problems where the search space can be characterized as an elementary landscape. For these search spaces the evaluation function is an eigenfunction of the Laplacian matrix that describes the neighborhood structure of the search space. Problems such as the Traveling Salesman Problem (TSP) and Graph Coloring are elementary.

Problems such as MAX-kSAT are a superposition of k elementary landscapes.

This research has exploited statistical and mathematical properties of elementary landscapes to develop new gradient methods for combinatorial optimization problems, particular for TSP and MAXSAT and other k-bounded pseudo-Boolean optimization problems.

**Distribution Statement****This is block 12 on the SF298 form.**

Distribution A - Approved for Public Release

**Explanation for Distribution Statement****If this is not approved for public release, please provide a short explanation. E.g., contains proprietary information.****SF298 Form**

DISTRIBUTION A: Distribution approved for public release.



Please attach your [SF298](#) form. A blank SF298 can be found [here](#). Please do not password protect or secure the PDF. The maximum file size for an SF298 is 50MB.

[AFD-070820-035.pdf](#)

**Upload the Report Document.** File must be a PDF. Please do not password protect or secure the PDF. The maximum file size for the Report Document is 50MB.

[finalreport15AH.pdf](#)

**Upload a Report Document, if any.** The maximum file size for the Report Document is 50MB.

**Archival Publications (published) during reporting period:**

33 archival publications are listed in the report.

An abstract is given for each publication.

**Changes in research objectives (if any):**

The project explored the TSP. Work on vehicle routing was not carried out, but extensive work was done on MAXSAT, pseudo-Boolean optimization, frequency assignment and graph coloring and related scheduling problems. The mathematics were a better fit for these problems. The extent of the work was well beyond our original expectations.

**Change in AFOSR Program Manager, if any:**

Don Hearn was the original Program Manager. The program is now managed by Fariba Fahroo.

**Extensions granted or milestones slipped, if any:**

There was a one year no cost extension.

**AFOSR LRIR Number**

**LRIR Title**

**Reporting Period**

**Laboratory Task Manager**

**Program Officer**

**Research Objectives**

**Technical Summary**

**Funding Summary by Cost Category (by FY, \$K)**

	Starting FY	FY+1	FY+2
Salary			
Equipment/Facilities			
Supplies			
Total			

**Report Document**

**Report Document - Text Analysis**

**Report Document - Text Analysis**

**Appendix Documents**

**2. Thank You**

**E-mail user**

Aug 31, 2015 19:03:46 Success: Email Sent to: whitley@cs.colostate.edu